A Bayesian Network for Combat Identification

Ing Sicco Pier van Gosliga, MSc
Ir Henk Jansen
TNO Physics and Electronics Laboratory
P.O. Box 96864
2509 JG The Hague
THE NETHERLANDS

E-Mail: vangosliga@fel.tno.nl / h.jansen@fel.tno.nl

SUMMARY

This paper reports the results of an investigation at TNO Physics and Electronics Laboratory for the Royal Dutch Navy on how a Bayesian network can be used to introduce more transparency in decision aid as to the level of confidence of the information that is used. To assess the feasibility of a Bayesian approach to new decision support system concepts, we have chosen to use the identification process of air contacts aboard navy frigates as a case study. The identification process is a time intensive and mind consuming process that is critical to anticipate an air attack. Wrong decisions may have fatal consequences and for example: identifying a neutral aircraft as an hostile aircraft may cost the lives of many people and may cause undesired political instability. Vice versa, mistakenly taking a hostile aircraft for a friendly one will give opponents the tactical advantages of a surprise attack. Typically a list of identification criteria is used by air defense personnel to discriminate hostile aircraft from neutral and friendly aircraft. These predefined criteria may change from mission to mission but will always follow a strict scheme of "if-then-else" clauses. Although identification involves reasoning with uncertain information, current procedures do not make this uncertainty explicit. Embedding Bayesian inference techniques in decision support systems would enable us to reason with uncertainty in a scientifically sound and consistent manner. Bayesian networks can express the likelihood of a hypothesis such as the identity of air contact being hostile as an explicit value, even when information about a contact is uncertain and incomplete. Making this uncertainty explicit, enables navy personnel to know how much confidence it should have the probability of hypotheses that are based on it.

BAYES THEOREM

Aircraft contacts, defined by tracks, are processed by automated information systems and presented to navy personnel that evaluate each track's identity and threat by interpreting its observed characteristics. Typical characteristics are [1][2]: altitude, speed and maximum speed, sudden manoeuvres, flight in formation, country of origin, adherence to air lane and adherence to air traffic control orders (ACO). There are multiple methods in use to determine the most likely identity for a track: electronic support measures (ESM), identification friend or foe (IFF), cross-told identification and visual identification. Each of these methods will deliver an identity for the track. Obviously air defence personnel are facing a problem when two methods suggest a different identity to a track. In a broader context this problem is commonly referred to as the Information Fusion problem. An approach to deal with such inconsistencies is using Bayesian statistics to derive the most likely identity.

The Bayesian theorem is based on the assumption that the probability of an event is dependent on the presence of other events, if causal relations exist between them. The basic mathematical rules for probabilistic reasoning are the product rule, the sum rule, the theorem of total probability and Bayes

Paper presented at the RTO IST Symposium on "Military Data and Information Fusion", held in Prague, Czech Republic, 20-22 October 2003, and published in RTO-MP-IST-040.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
00 MAR 2004	**N/A**	**-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
A Bayesian Network for Combat Identification	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **TNO Physics and Electronics Laboratory P.O. Box 96864 2509 JG The Hague THE NETHERLANDS**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM001673, RTO-MP-IST-040, Military Data and Information Fusion (La fusion des informations et de données militaires)., The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**UU**	**19**	

theorem (where Bayes rule is actually a consequence of the product rule). Since especially the last rule forms the foundation of all Bayesian systems it will be discussed briefly. We will use 'P(h | e)' as a formal notation for the probability (i.e., the degree of belief) of hypothesis 'h' being true, given circumstantial evidence 'e'. The probability of an event is expressed as a value scaling from 0 to 1, where the value 1 (or 100%) says that the event certainly takes place or is believed to take place and 0 (or 0%) definitely not.

$$p(h \mid e) \times p(e) = p(e \mid h) \times p(h)$$

Figure 1: Bayes theorem.

The left-hand term, 'P(h | e)', is called the posterior probability. It represents the probability of hypothesis 'h' after considering the effect of evidence 'e'. The second term, 'p(e)', returns the chance of observing 'e'. The term 'p(e | h)' is called the likelihood, and it stands for the probability of circumstantial evidence 'e' assuming the hypothesis 'h' is true. The last term, 'p(h)', is the prior probability of event 'h' reflecting our belief in 'h' without any additional information (i.e. without evidence 'e').

$$p(h \mid e_1 , e_2) = \frac{p(h \mid e_2) \times p(e_1 \mid h , e_2)}{p(e_1 \mid e_2)}$$

Figure 2: An alternative (extended) notation of Bayes theorem, where 'e_2' is additional evidence.

This paper discusses how to embed domain knowledge into a Bayesian network. A Bayesian network is a graphical model that represents the joint probability distribution for a large set of variables. Graphical models are graphs in which nodes represent variables, and the arcs represent the relations between them. We make distinction between two types of variables: observed variables and hypotheses. Being a graph structure, a Bayesian network can be manipulated with a large variety of mathematical methods from graph theory. Furthermore several belief-updating algorithms based on the Bayesian theorem, have been developed for such networks. Hence, the name: Bayesian belief network. Given that algorithms based on the Bayesian principle are mathematically consistent [3] (as opposed to e.g. Fuzzy Logic, Dempster-Shafer and uncertainty factors) we have chosen to use a Bayesian network structure to approach the combat identification problem.

Observed variables are those variables that have been assigned a value. The hypotheses are variables whose values are derived from the observed variables, but of which the actual values are unknown. For hypotheses an approximated value can be derived given the values of observed variables (e.g. 'max[p(h | e_1 , e_2)]'), based on a priori knowledge (e.g. 'p(e_1 | h , e_2)'). This paper will not discuss how the Bayesian theorem is used to derive this information. People who are interested in the underlying algorithms be may be interested in the following literature: [4] [5]. The focus of this paper lies on how a domain can be modelled into a Bayesian network and uses the combat identification process as an example. Designing a Bayesian network takes three phases, which will be described now.

PHASE 1: DEFINING VARIABLES AS HYPOTHESES

Since the nodes in a Bayesian network form the basis of the graph, we need to define the variables first. Typically a set of variables is needed that represent facts that can be observed. The variables that cannot be observed will be treated as hypotheses. Common criteria for identifying tracks have been mentioned above, before we discuss how these criteria relate to each other we will discus how they can be used as variables in a Bayesian network model. The variables can be interpreted as representing so-called events. A variable may in theory have any number of states. A variable may, for example, be the speed of an

aircraft (300 m/h, Mach 2 etc.) or the identity of a track (neutral, hostile or friendly). Important is that a variable may only be in exactly one state at a time. An aircraft cannot fly 300m/h and 600 m/h at the same time; neither can a track be hostile and neutral at the same time. The state may be unknown to us. When the state of a variable is unknown, the known states of other variables and prior knowledge can give us an indication of what the unknown state may be.

The variable "track identity" is the central hypotheses. In the case discussed in this paper, its value is unknown. To assess the most likely state of "track identity" observations are necessary. Prior knowledge alone is inadequate to reach the required level of certainty. Such prior knowledge may be the known ratio of friendly planes to neutral and hostile planes in our neighbourhood. This prior knowledge is as dynamic as the surrounding world is. The mentioned criteria that are used for combat identification will be embedded as variables in our model.

On decision level, the set of feasible actions is limited and therefore discrete. For this reason we will strictly use discrete variables. Continuous variables (e.g. speed and altitude) will be discretized. There are also practical reasons why discrete variables are preferred. To support the decision-making process the likelihood of each state must be calculated for all unobserved variables. The total probability of all these states should sum up to 1. When the set of values for a variable would be very large, the probability for a value will be close to zero. Such small numbers introduce numerical problems and unnecessarily increase the computational complexity of the problem. The smaller the range, the quicker the algorithm is. There are Bayesian algorithms that can be used on continuous variables. However these algorithms use restrictions on the format of probability density functions (e.g. Gaussian distribution functions) or do not use a complete search of the solution space (e.g. Monte Carlo simulation). People who are interested in using continuous variables in Bayesian network are advised to read the following literature: [6].

$$\text{speed} \in \{0..k, k..2k, \ldots, s_{max} - k \ldots s_{max}\}$$

$$\text{altitude} \in \{0..a, a..b, b..c\}$$

Figure 3: Instead of continuous, discretized versions will be used for speed and altitude.

If we want to use continuous variables, they should be converted to discrete variables. To illustrate this process, we will discuss the variable speed that we will use for the speed of a track. First of all, we need to limit the range of this variable. We assume that a track will not fly faster then a certain speed (say s_{max}) and we won't accept speeds below zero: $0 < \text{speed} < \text{smax}$. We can split this range in a number of equal sized parts, e.g. $\text{speed} \in \{0..k, k..2_k, \ldots, s_{max} - k \ldots s_{max}\}$. This gives us a discrete version of the continuous variable speed. Obviously the size of the set affects the accuracy of the model. Rather than using a high resolution, it is more effective to choose a range for speed that discriminates various possible identities of a track. Since civil planes in general fly within a certain speed range. Military jets may exceed the speed of civil planes, but not necessarily (e.g. military helicopters fly at lower speeds). A similar process is used for the variable altitude.

$$\text{adherence\_to\_ACO} \in \{\text{true, false}\}$$
$$\text{adherence\_to\_airlanes} \in \{\text{true, false}\}$$
$$\text{sudden\_manoeuvres} \in \{\text{true, false}\}$$
$$\text{flight\_in\_formation} \in \{\text{true, false}\}$$
$$\text{identification\_manoeuvre} \in \{\text{correct, incorrect}\}$$

Figure 4: The variables that will be used to describe a track's behaviour.

Commercial planes are obliged to follow civil air lanes. These lanes are static and commonly known. The recorded history of a track can be used to determine whether the plane was adherent to an air lane.

When an airliner has been outside its air lane, it might be an airliner in trouble or a hostile aircraft in disguise. Whether a plane is coherent to an air lane can be expressed as a fact (true or false), alternatively the distance to the mean of an air lane can be used to express the degree of belief of coherence. The same approach can be used for the other behavioural characteristics: adherence to air traffic control orders, sudden manoeuvres and flight in formation. An instructed manoeuvre is performed by friendly forces for identification purposes. Since new instructions are given on a daily basis and are restricted knowledge, performing a correct manoeuvre is considered to be a reasonable hard criterion to positively identify friendly forces. We will only use "correct" and "incorrect" in our model.

ID_by_ESM ∈ {neutral, hostile, friendly}
IFF_mode∈{mode 3/ac, mode 4, other}
ID_by_visual ∈ {neutral, hostile, friendly}
ID_by_intel∈ {neutral, hostile, friendly}
ID_by_FAC∈ {neutral, hostile, friendly}
ID_by_ally∈ {neutral, hostile, friendly}

Figure 5: Various methods can be used to derive an identity for a track. However, the results of these methods may differ. This set of discrete variables will be used to hold their values.

Electronic warfare support measures (ESM) forms the division of electronic warfare involving action taken to search for, intercept, identify and locate radiated electromagnetic energy for the purpose of immediate threat recognition. It provides a source of information required for immediate decisions involving electronic countermeasures, electronic counter-countermeasures and other tactical actions such as avoidance, targeting and homing. ESM is limited as a method of identification since it relies on goniometry to relate sensed radiation to a specific track. Therefore its readings are less reliable in dense air traffic, because it will be uncertain which sensor reading relates to what track. One might consider adding an extra variable to describe the density of air traffic around the assessed track, when using EMS as a criterion for identification. Characteristic patterns of electromagnetic radiation from known planes can be measured effectively and used as a reference. The availability of usable reference material affects the reliability. Another automated way to determine a track's identity is using IFF, "identification friend or foe". This system was developed as a military system to discriminate hostile forces from friendly forces. All IFF systems use a "challenge and response" system. In civil systems the ATC Beacon system will "challenge" an airborne transponder. The challenge will be for a specific mode and wait for a response. The aircraft's transponder will determine which mode was used and reply with an appropriate response. Modes are determined by the timing between the challenge pulses. IFF has 4 modes:

- The first mode originates from the original military system that was used in military air traffic control to determine what type of aircraft is answering or what type of mission it is on.
- The second mode is a newer system that replaced the first, also only for military use. This mode requests the tail number to identify a particular aircraft.
- The third mode is in use for civilian systems. It knows optional additional modes:
  - Sub-mode A is used internationally to identify airliners for air traffic control.
  - Sub-mode C altitude encoding for air traffic control (normally used in conjunction with 3/A).
  - Sub-mode is WIC war identification code for military use.
- Finally mode 4 is the military system that is in use nowadays. Instead of modes 1 and 2, this mode is encrypted and therefore more secure. It uses NATO specific crypto-secure IFF broadcast messages and is considered to be the most reliable criteria to determine the identity of a track.

IFF is a discrete variable. The introduction of mode 3 and mode 4, made the first two modes obsolete. It is unlikely that they will be used in wartime. Airliners normally use mode 3 in conjunction with sub-modes

A and C. Air traffic control relies on airliners to use these modes, therefore we will only accept mode 3 when both sub-modes are used. This leaves us with NATO secure mode 4 and the value "other reply" for all replies that are not covered by both states. Note that the value "other reply" is not the same as "no reply observed". Normal visual information and visual images derived from the optical monitoring of the electromagnetic spectrum from ultraviolet through far infrared can be used effectively to assess the identity of tracks. The visual identification process is traditionally performed by human agents. Automated systems match observed images to data and patterns stored in a database. The reliability is dependent on the quality of the image and the visibility of distinguishing features. When another party (e.g. a forward air controller, an ally or intelligence) has identified a track, the reliability of this identity is dependent on the source. Since multiple sources may provide a cross-told identity for a single track. Therefore the values for information source are not mutually exclusive and a separate variable for each source required.

PHASE 2: DESIGNING A CAUSAL NETWORK

Bayesian networks or belief networks are directed graphical models. Directed graphical models take into account the directionality of the arcs that link the nodes together and are not allowed to have directed cycles. This has several advantages of which the most important is that one can regard an arc from A to B as indicating that event A "causes" event B. Thus the arc direction is determined by chronological order of events. Using causal relations as a guide to construct the graph structure, leads to a reasonably intuitive model of its domain. In addition causal structures can be learned by fitting the strength of the causal relation to data. This can be convenient in problem areas were domain modelling is problematic.

Although the direction of each causal relation must point into one single direction, it does not mean that the state of a child will not affect the likelihood of its parent's states. The likelihood of events that form the circumstantial evidence of hypotheses may on their turn depend on other criteria. For example: for a track the adherence to an air lane depends on its position and course, while the probability of that track having a certain position may depend on the weather conditions. It is forbidden to use bi-directional arcs. In practice this means that causal cycles like the chicken-and-eggs problem are not allowed. On occasions where the directionality of a causal relation is not clear it is hard to construct a network. It is not allowed to have $A \rightarrow B$ and $B \rightarrow A$ in the same network, because cycles are prohibited. Normally, $A \rightarrow B$ and $B \rightarrow A$ should be redundant to each other. Because Bayes product rule is symmetrical, $[P(A \mid B) \times (B) \Leftrightarrow P(B \mid A) \times P(A)]$, there should be no functional difference. If it is not possible to translate $A \rightarrow B$ into $B \rightarrow A$ and vice versa, the model is semantically incorrect. The problem is likely to be caused by the fact that two separate events are represented by one and the same variable: $A \rightarrow B$ and $B \rightarrow A'$, where A' is in fact a different instance than A. Another cause may be that A and A' are referring to the same instance, but at a different moment in time. Our model is centred on the identity of a track. In the previous section we introduced a set variables based on criteria that are normally used for combat identification. We will now bind these variables together in a network structure based on the causal relations between them. Some of the causal relations are strengthened in the presence of other events. For these events extra variables will be introduced.

Figure 6: On the left a decision tree is shown for the identity of a track. On the right a graph is seen that shows the structure of the causal relations between altitude, speed and identity.

Both models in figure 6 describe the relation between altitude, speed and identity. The left model describes how the identity of a track is *determined* by altitude and speed, where the right model describes how the speed and altitude are *influenced* by identity. What is the exact difference? If we consider all nodes as variables describing events, the timing of these events defines the direction of the arc. When a track has a certain identity, say hostile, it is likely to fly low at a high speed. The track was already hostile before it decided to fly low. It did not become hostile because it was flying at a high speed or a low altitude. A counterexample: a civil plane can fly low, when it is nearing an airport. The airport did not suddenly appear because the plane was flying low; of course it was already there. Since the airport is near, the pilot is triggered to fly low.

identity $\Rightarrow$ altitude $\wedge$ speed
near airport $\Rightarrow$ altitude
near airport $\in$ {near, not present}

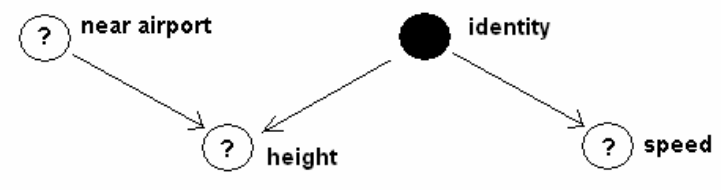

Figure 7: The graph on the right shows how knowing the state of 'identity' blocks the relation between 'speed' and 'altitude'. On left a textual representation of the same causal relations.

The direction of the arcs guides the belief update algorithm through the network. One way of looking at a Bayesian network is that observed events trigger other events that are directly related to it to update the likelihood of their states. These updated events will trigger related events on their behalf, and so on. In that way belief updates will propagate though network away from the observed events. This is why the direction of the arcs is that important. Cycles would create an endless loop for such belief updates, exactly as it does in the chicken-and-egg problem. The directionality guarantees that the belief update will stop at certain points in the network. Without going into further detail on how Bayesian inference algorithms work, the small network in figure 7 illustrates how belief updates are "blocked" by directed arcs. The figure illustrates the following relations: "a hostile aircraft is likely to fly at a low altitude to avoid radar contact" and "a civil plain is not likely to fly that low, unless it is landing". Suppose we only know the *altitude* of a plane, then the *altitude* gives us a clue of the *identity* of a plane, hence the most probable *speed*. When we do know the *identity*, but do not know the altitude. Knowing the *speed* will not help us to estimate the *altitude*, since there is no direct relation between them, only via *identity*. The knowledge of *identity* blocks the relation. Knowing identity dictates the likelihood of *altitude*. Only knowing whether the near is near an airport force us to reassess the probabilities of *altitude*. If *track near airport* is unknown, the prior probability of *track near airport* is used to derive the most likely state of *altitude*.

identity $\Rightarrow$ made sudden-manoeuvres
heavy turbulence $\Rightarrow$ made sudden-manoeuvre
identity $\Rightarrow$ flight in formation
identity $\Rightarrow$ adherence to air lane
turbulence $\in$ {true, false}

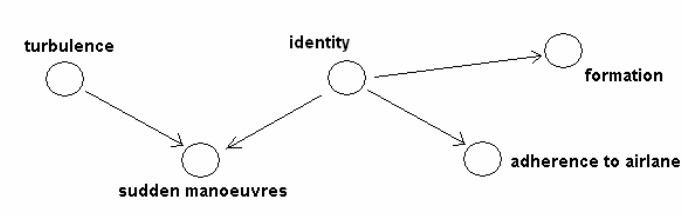

Figure 8: Graph structure of the causal relations between 'identity' and flight characteristics.

A track's identity can be recognised by several behaviours that are typical for civil or military aircraft. Civil tracks follow flight lanes and keep a steady course and speed. Heavy turbulence may cause sudden manoeuvres as well, but will naturally not affect the identity of a track. Military tracks normally fly in

formation and may use sudden course changes on their path. Carriers are unable to make fighter-jet manoeuvres.

$$\text{identity} \Rightarrow \text{ESM}$$
$$\text{identity} \Rightarrow \text{IFF}$$
$$\text{tracks close by} \Rightarrow \text{IFF}$$
$$\text{tracks close by} \Rightarrow \text{ESM}$$
$$\text{tracks close by} \in \{\text{true, false}\}$$

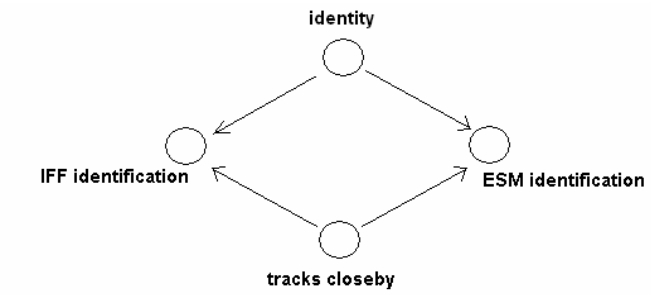

Figure 9: On the right a graph structure showing the causal relation between identity and identification methods IFF and ESM. Note the double dependencies.

Various technical aids are available to discriminate tracks. Established methods are Electronic Support Measures and the NATO secure IFF protocol. ESM identifies tracks by using pattern matching techniques on electromagnetic energy emitted by the track. Radiated energy is related to known characteristics. ESM is less reliable when multiple tracks are flying close to each other, because on such occasions it will be hard to pinpoint the source of the radiation. IFF has this same problem. Figure 9 shows the causal relations between these methods and a track's identity. Together with the three causal relations for the cross-told identities (allies, intelligence and forward air controllers) the relations in this section make up a Bayesian network as shown in figure 10. Note that this network does not represent a decision tree of the identification process. In a decision tree, the arcs would be the other way around. The network is a causal model centred on the identity of a track, describing a variety of factors that are symptoms of it.

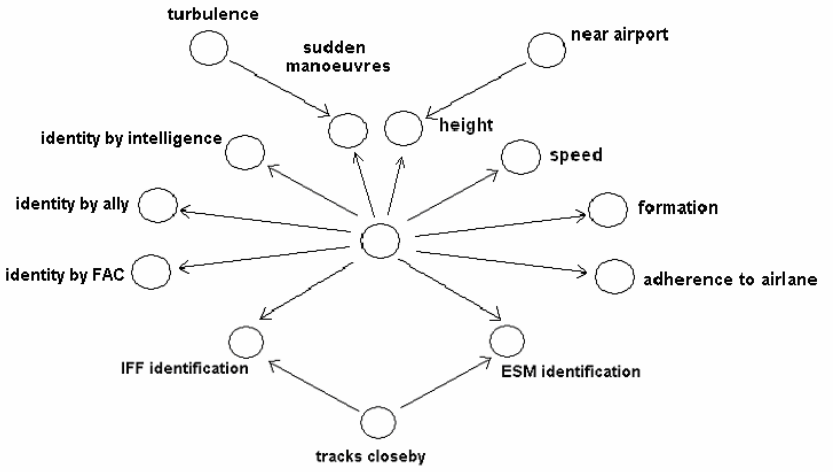

Figure 10: This graph structure gives an overview of all causal relations described in this paper.

PHASE 3: DEFINING THE PRIOR PROBABILITIES

Bayesian inference relies on probability distributions. Defining these distribution is a challenging process, since many probabilities are not known precisely, if not impossible to know exactly. A probability function describes the probability of a variable having a certain value, given the values of other variables to which a causal relation exists. A probability function over discrete variables can be structured as a

probability table. In the previous section the following variables and correlations were introduced to describe the relations between the identity, speed and altitude of an aircraft.

altitude ∈ {0..a, a..b, b..c}
identity ∈ {hostile, friend, neutral}
speed∈ {0..k, k..2k,… ,s_{max} - k …s_{max}}
near airport ∈ {near, not present}
identity ⇒ altitude
identity ⇒ speed
near airport ⇒ altitude

p(airport = present) = 50%
p(airport = not present) = 50%
p(identity = hostile) = $33^1/_3$%
p(identity = neutral) = $33^1/_3$%
p(identity = friendly) = $33^1/_3$%

Figure 11: On the left a set of variables is seen and how they relate to each other. Two variables have no dependency on another variable (i.e. not present on the right-hand side of a relation). For both variables, the prior-probability is show on the right.

Table 1: this table holds all probabilities for the function P(altitude | identity, airport)

		altitude		
identity	**airport**	0..a	a..b	b..c
neutral	not present	10%	80%	10%
neutral	present	70%	20%	20%
hostile	not present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%
hostile	present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%
friendly	not present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%
friendly	present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%

For the part of network that encapsulates these variables (seen in figure 7) the following probability functions need to be defined: '*P(altitude | identity, airport)*', '*P(speed | identity)*', '*P(airport) and P(identity)*'. That is one table for each variable. The number of relations that affect a variable defines the number of dimensions for the table that describes its prior probabilities. The variables *airport near* and *identity* are independent of other events. If the likelihood of all states of a variable *v* are equal ($p(v=s_1) == p(v=s_2) == p(v=s_3)$ *etc.*), this means that the prior probabilities of *v* will not play a discriminating role of influencing the probabilities of the states of variables related to it. Basically this means that we do not want to assume anything on the presence of an airport when we do not know whether this is the case. Neither we want to assume anything on the identity. We can safely assume that the position of a track is always known. Given a database of all airports we could use the position of the track to determine whether an airport is near or not. In that case *position* is an observed variable, the state of variable *airport* will not have to be estimated. We could also use the position of a track to say something of the likelihood of a track having a certain identity. When we know how many friendly aircraft are in the area, we may have an estimate of the number of neutral and hostile planes in the area as well. This information could be used to dynamically generate prior probabilities, based on actual knowledge of the neighbourhood. The variable identity will remain unobserved, but if nothing is known about the context of a track except its location we could say something about the probability of that track being hostile. Naturally since the numbers of friendly in the area shifts over time, so do the probability tables of *identity*.

A method to determine the likelihood of each value for each context is learning from recorded data. Statistic methods can be used to estimate the chance of value based on the frequency of occurrence in the data for a certain context [7]. For many applications it is unlikely that training can learn all possible scenarios, because all possible states of the context will have to be available in the data. This is especially

the case for combat identification. Even if this requirement is met, it is still not certain that recorded data is representative for future missions because of the ever shifting paradigm in dynamic environments. It is possible however to train subparts of the network which function within a relatively static system.

BENEFITS

By recognizing causal relations between criteria may be weak or strong depending on their context, it is better to include this knowledge in the decision-making process. Various tools can aid this process, but not all of them will take into account the dynamics of causality. Probabilistic models are an attempt to model causality in a less rigid way than rule-based systems. A Bayesian-based system can evaluate the identity of a track based on incomplete information, because it is able to estimate the values of the unobserved variables by using the prior knowledge and relate this to known observations. Not relying on complete information makes the system more reliant. When not all the information can be available at once, the system can determine the probability of posterior events and anticipate on the most likely before it has complete knowledge of what is to come. For the identification case this means, that all tracks can be given a likely identity even if certain information such a IFF or ESM is not yet available. Bayesian networks have the benefit to be able to determine the priority of missing information given the current context, without being vulnerable to tunnel vision. Because of this the information gathering process can be directed efficiently. Furthermore, a Bayesian-based system knows the weight (relevance) of all used criteria. For each hypothesis the most important reason can be given why it is valid or not. Rather than listing all reasons, it can highlight the most important one for the current situation. In a time-critical decision-making process, information overflow is a widely addressed problem. Context aware information filters have been proposed as a solution to ease the stress on decision makers. A Bayesian-based system can be used as a context aware information filter. The ability of Bayesian methods to prioritise information on relevance makes an effective tool for information presentation and gathering.

Knowledge in a Bayesian network is scalable. Adding new variables does not affect the prior knowledge. Adding causal relations, only affects the prior knowledge of the variables that are directly related to it. In a rule-based system the set of rules affects the outcome of the inference mechanism. For some mechanism also the order in which the rules are addressed is of importance. If rules are added to the system, one has to be aware that each extra rule can alter the functioning of the whole rule-set. To prevent undesirable effects, complete knowledge of the rule-set is necessary. Because of this, maintenance on rule-based system is a elaborate process and will never be a routine operation. Compared to rule-based systems, Bayesian-based systems can be scaled up easier.

CONCERNS

Probabilistic reasoning itself will not be the main challenge for the technical implementation of a Bayesian-based decision support system. Algorithms have been extensively described in scientific papers and implementations of these algorithms are available as ‘of-the-shelf’ software components from a broad range of parties. Most effort will be required for the construction of a sound model and the human factors of reasoning with uncertainties, presenting the output of the reasoning algorithm.

Special interest should be given to modelling and maintaining the knowledge embedded in the network. Since Bayesian systems use a model-based approach, a sound understanding of the domain is required. In the past expert systems have been usually based on heuristic knowledge. For model-based systems it is favourable not to use heuristic rules, when they are based on a weak correlation and not founded on understanding of the underlying mechanisms. Typically heuristics are simple diagnostic rules that draw a strict conclusion based on observed symptoms, while these symptoms may not be directly related to the cause. Such rules are out of place in a causal network, however rules that are founded on causal relations can be translated and will fit in [8]. In this paper the variables are presented as events. Although an event

is related to a moment in time, this property is not exploited in the architecture of the described causal model. The presented model takes only into account the *occurrence* of events. However, knowing *when* and in *what order* events occur gives us a strong indication on the cause of events. Research in the field of probabilistic modelling and reasoning currently concentrates on embedding the dynamic aspects of time into Bayesian belief networks. Although these aspects have not been discussed in this paper, the authors would strongly advice engineers involved in causal modelling to be aware of the crucial role of time and make use of the latest developments in the field of dynamic Bayesian networks.

In this paper probabilities are represented as percentages. Research exposes that people find it hard to interpret hypotheses in such a notation [9][10]. Care should be given on how to represent probabilities. An alternative notation is the frequency notation, which expresses the probability of an event as the frequency of occurrence of one single event. A frequency-based notation seems to work better then percentages, but is not applicable for all variables and domains. Suppose a frequency based notation is used for the Air Defence case, how should someone act on a computer generated sentence: "On 1 in 3 times under these circumstances that plane would be hostile."? A graphical representation avoids the textual interpretation problems. For the combat identification case a triangular representation was used (as seen in figure 12). An equilateral triangle can visualize three related dimensions such as hostile, friendly and neutral. Each corner of the triangle represents the upper boundary of a dimension. Because the values of a probability variable are mutually exclusive, a triangle can visualize each event that has three possible states. In this way hostile targets will be positioned in another corner then the neutral tracks. The tracks that have an unknown identity (i.e. not enough is known to discriminate between them) will be placed in the middle of the triangle. One advantage of this visualization is that hostile tracks are visually discriminated from other tracks, and found on a fixed location. For displaying certainties on a topological display, colour codes can be used to add an extra layer of information on a map. When all possible states of an event are represented by a colour; the mixture of these colours represents the identity of a track. This method may have some disadvantages, like incompatibility with the standard NATO colour codes and symbols. Although some possible ways of visualizing uncertainty have been examined, it has not been the focus of our research at TNO Physics and Electronics Laboratory. However, since research indicates that people have difficulties in both reasoning with and interpreting uncertainties [10], it is important for the acceptance of decision aid systems that human factors are given full attention in the engineering phase. The authors encourage research in this field.

Finally the problem remains what to do when a track seems to be hostile. This however is not mainly – or at all– an engineering issue but a military and political one. When is a track hostile enough to act? Release criteria have been unaddressed in this paper, although they are closely related and triggered by identification criteria. The benefit of using decision aid is to support the decision makers by delivering plausible and solid conclusions on the identity of each track during the dynamic process of air combat. By making uncertainty explicit, de decision process is getting more transparent.

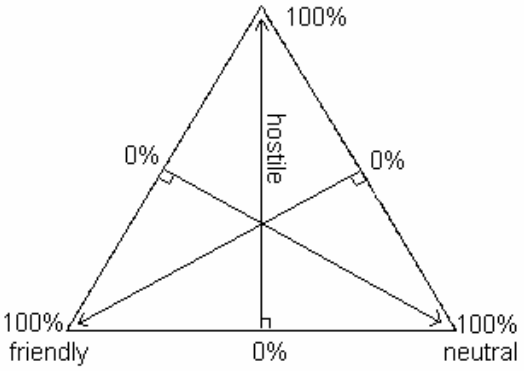

Figure 12: A triangular representation of a tracks identity.

CONCLUSION

By incorporating Bayesian techniques into the decision loop to assist in the identification process, time and energy can be saved. The model presented in this paper is merely a sketch of the actual Bayesian model for air defence identification and serves merely to illustrate our point of view on how such a model should be designed. The availability of software implementations of multiple Bayesian algorithms makes the technical implementation relatively easy. In general, designing an adequate and sound causal model of application domain will be the greatest challenge. Intimate knowledge of Bayesian statistics will not be a necessity, but planners do require a thorough understanding of the subject. Concerns remain however on how to present the results to navy personal and how the knowledge in the network should be kept up to date. The authors expect that further research may result in a better way to incorporate time aspects in Bayesian models and better ways to present the outcomes to navy personnel.

REFERENCES

[1] NATO Military Agency for Standardisation, "STANAG 4162 Technical Characteristics of the NIS Identification Data Combining Process", March 2001 (NATO RESTRICTED)

[2] NATO, Military Agency for Standardisation, "STANAG 4162 Annex A: Source type description, the standardized requirement", March 2001 (NATO RESTRICTED)

[3] C. P. Robert, Springer "The Bayesian Choice", 11:507-418, 2001

[4] N. J. Nilsson, Morgan Kaufmann "Artificial Intelligence, A New Synthesis", 19:317-342, 1998

[5] F. V. Jensen, Springer Verlag "Bayesian Networks and Decision Graphs (Statistics for Engineering and Information Science)", 2001

[6] R. Schachter, C. Kenley, Management Science "Gaussian Influence Diagrams", 35:527-550, 1989

[7] D. Heckerman, Microsoft Research Technical Report "A Tutorial on Learning with Bayesian Networks", March 1995

[8] M. Korver, P.J.F. Lucas, Medical Informatics 18(3): "Converting a Rule-Based Expert System into a Belief Network", 1993

[9] G. Gigerenzer, P. Sedlmeier, Journal of Experimental Psychology "Teaching Bayesian Reasoning in less than two Hours", September 2001

[10] G. Gigerenzer, U. Hoffrage, Psychological Review "How to improve Bayesian Reasoning without instruction: Frequency Formats", September 2001

NATO RTO Symposium on "Military Data and Information Fusion"

A Bayesian Network for Combat Identification

TNO Physics and Electronics Laboratory

Combat Identification

- TNO Physics and Electronics Laboratory
- Royal Dutch Navy
- making uncertainty explicit
- more transparency
- knowing the level of confidence

Bayes Theorem

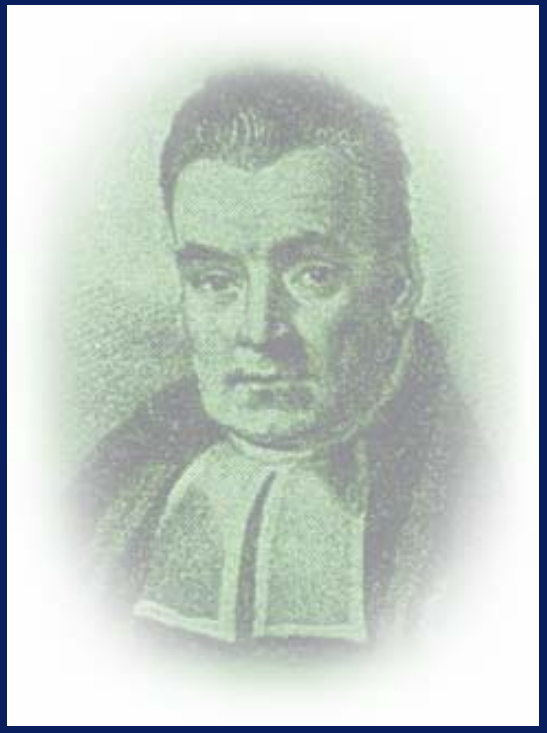

$$p(h \mid e) = \frac{p(e \mid h) \times p(h)}{p(e)}$$

I: Variables for Events

speed $\in$ {0..k, k..2k,… ,s_{max} - k …s_{max}}

altitude $\in$ {0..a, a..b, b..c}

adherence_to_ACO $\in$ {true, false}

adherence_to_airlanes $\in$ {true, false}

sudden_manoeuvres $\in$ {true, false}

flight_in_formation$\in$ {true, false}

identification_manoeuvre $\in$ {correct, incorrect}

ID_by_ESM $\in$ {neutral, hostile, friendly}

IFF_mode $\in$ {mode 3/ac, mode 4, other}

ID_by_visual $\in$ {neutral, hostile, friendly}

ID_by_FAC $\in$ {neutral, hostile, friendly}

ID_by_ally$\in$ {neutral, hostile, friendly}

II: Causal Relations

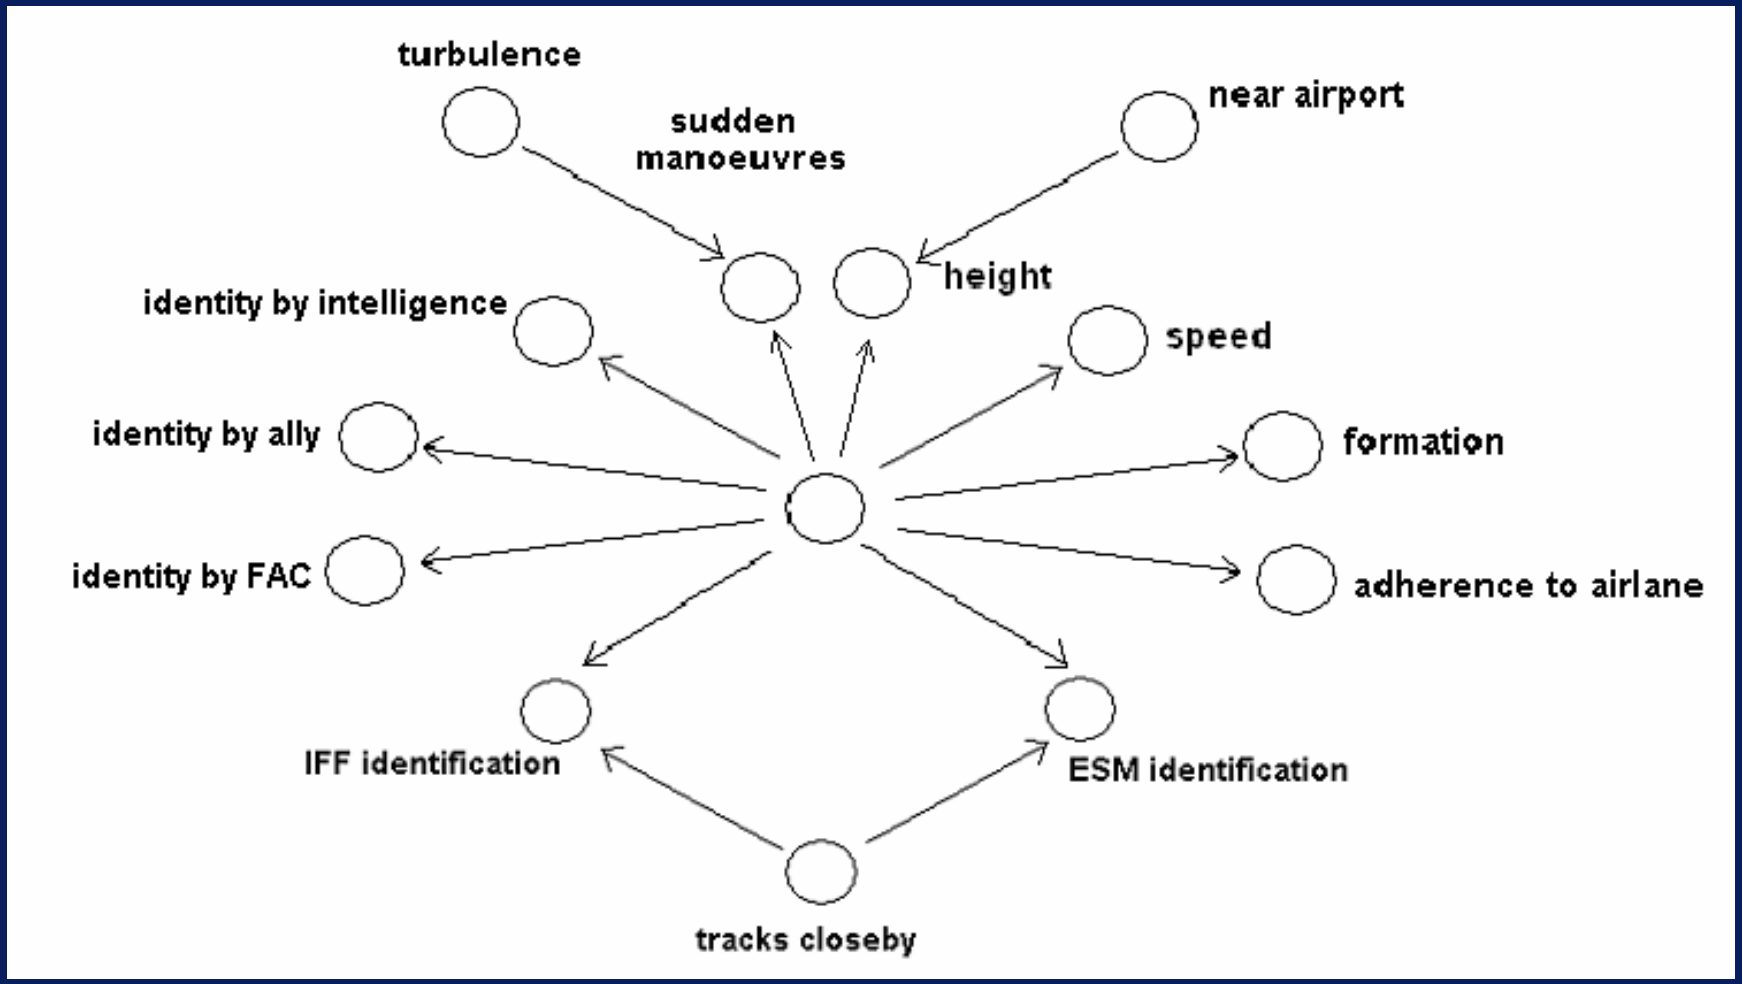

III: Probability Distributions

altitude $\in$ {0..a, a..b, b..c}
identity $\in$ {hostile, friend, neutral}
speed$\in$ {0..k, k..2k,… ,s_{max} - k …s_{max}}
near airport $\in$ {near, not present}

p(airport = present) = 50%
p(airport = not present) = 50%
p(identity = hostile) = $33^1/_3$%
p(identity = neutral) = $33^1/_3$%
p(identity = friendly) = $33^1/_3$%

P(altitude | identity, airport)

identity	**airport**	**altitude** 0..a	a..b	b..c
neutral	not present	10%	80%	10%
neutral	present	70%	20%	20%
hostile	not present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%
hostile	present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%
friendly	not present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%
friendly	present	$33^1/_3$%	$33^1/_3$%	$33^1/_3$%

Conclusion

Concerns

- world models must be kept up to date with reality
- humans have difficulties with interpreting uncertainty

Benefits

- scalable model
- standard framework
- information overflow
- consistent way of dealing with uncertainty, with a solid scientific foundation